

SENSOR WEBS: AUTONOMOUS RAPID RESPONSE TO MONITOR TRANSIENT SCIENCE EVENTS

Dan Mandl, Sandra Grosvenor, NASA Goddard Space Flight Center, Greenbelt, MD
Stu Frye, Mitretek Systems, Inc., Falls Church, VA

Robert Sherwood, Steve Chien, Ashley Davies, & Ben Cichy, Jet Propulsion Laboratory, Pasadena, CA
Mary Ann Ingram, Georgia Institute of Technology, Atlanta, GA

John Langley, Saquish Group, Half Moon Bay, CA

Felix Miranda, Richard Q. Lee, Robert Romanofsky, & Afroz Zaman, Glenn Research Center, Cleveland OH
Zoya Popovic, University of Colorado, Boulder, CO

ABSTRACT

To better understand how physical phenomena, such as volcanic eruptions, evolve over time, multiple sensor observations over the duration of the event are required. Using sensor web approaches that integrate original detections by in-situ sensors and global-coverage, lower-resolution, on-orbit assets with automated rapid response observations from high resolution sensors, more observations of significant events can be made with increased temporal, spatial, and spectral resolution. This paper describes experiments using Earth Observing 1 (EO-1) along with other space and ground assets to implement progressive mission autonomy to identify, locate and image with high resolution instruments phenomena such as wildfires, volcanoes, floods and ice breakup. The software that plans, schedules and controls the various satellite assets are used to form ad hoc constellations which enable collaborative autonomous image collections triggered by transient phenomena. This software is both flight and ground based and works in concert to run all of the required assets cohesively and includes software that is model-based, artificial intelligence software. Further more, experiments in more cost-effective interconnectivity between the various satellites are being conducted using adaptive antenna arrays in an architecture similar to cell phone towers. These experiments are being conducted now by a team of researchers and scientists at NASA's Goddard Space Flight Center and Jet Propulsion Laboratory in collaboration with numerous universities and other government agencies under the mantle of the Sensor Web. This activity provides a true end-to-end approach for prototyping the "system of systems" needed for global Earth observations as well as for honing lunar/planetary exploration strategies.

Keywords: autonomous detection, sensor webs, ad hoc constellations, progressive mission autonomy, collaborative remote sensing, smart antennas, adaptive antenna arrays

1. INTRODUCTION

NASA's Earth Observing 1 (EO-1)¹ spacecraft has been used and continues to be used as the core satellite to demonstrate various mission autonomy technology. The

original mission of EO-1 was to validate a number of space technologies during its first year. After successful completion of the first year objectives, the mission evolved into an on-orbit testbed for sensor web concepts. This paper will outline the overall operations concept for the sensor web and highlight some details of both the research being conducted in progressive mission autonomy and touch on related work in ground adaptive antenna arrays for low earth orbiting satellites, performed with EO-1, that could further facilitate future sensor webs.

The concept for sensor web is to link together ground and space-based instruments with event driven detections of scientific interest through a seamless set of software and communications interactions that weave together observation campaigns in an ad hoc sensor constellation and supply multiple data acquisitions as rapidly as possible and in as much depth as possible in a given time period for the purpose of autonomously capturing and categorizing transient phenomena.

The next-generation science and exploration systems will employ new observation strategies that will use multiple sensors in a dynamic environment to provide high quality monitoring, self-consistent analyses and informed decision making. The sensor web experiments described herein provide a prototype to explore the nature of automation necessary to enable dynamic observing of Earth and other planetary phenomena. The tools being developed improve our ability to autonomously monitor multiple independent sensors and coordinate reactions to better observe the dynamic phenomena. These systems enable users to specify events of interest and how to react when an event is detected. The systems monitor streams of data to identify occurrences of the key events previously specified by the scientist/user. When an event occurs, the system autonomously coordinates the execution of the user-desired reactions between different sensors. The information can be used to rapidly respond to a variety of fast temporal events without human intervention.

Many geophysical phenomena are dynamic and coupled. In order to fully understand them, we need to monitor them and obtain timely coordinated multi-sensor observations from widely dispersed instruments. The need for dynamic coordinated multi-sensor observations has given rise to the concept of sensor webs, which characterize future observing systems concepts more capable than today's independent observing systems.

^{*}Corresponding Author Address: Stuart Frye, Mitretek Systems, 3150 Fairview Park Drive M/S 540, Falls Church, VA, 22042-4519. Stuart.W.Frye.1@gsfc.nasa.gov

Sensor webs will monitor the intrinsically dynamic behavior of a wide variety of naturally occurring (e.g., wild fires, flash floods, hurricanes, volcanoes) and human-induced (e.g., toxic spills, pollution) events and phenomena. It is envisioned that sensor webs will:

- Maximize the return of only the most useful scientific information;
- Minimize overall response time of the system when monitoring rapidly evolving phenomena;
- Conduct near-real-time synthesis or "fusion" of information from multiple assets. Numerical models will also be considered as assets and will contribute toward improving our understanding of the complex, interrelated processes that drive the formation and evolution of environmental phenomena.

Unlike today's "stove-piped" missions where scientific goals are often achieved using passive observing strategies, the dynamic nature of the future observing systems implies that these goals can be achieved using reactive and proactive dynamic observational strategies that use all available interconnected resources (e.g. various spacecraft, balloons, numerical models and other distributed/configurable sensors). Our experiments begin to investigate the potential benefits of sensor webs to determine the properties that they should possess, and to determine the best methods to develop, integrate, deploy and operate them. Furthermore, we investigated some of the needed technologies and information systems needed to enable sensor webs.

It was noted that as more assets are placed in orbit, more opportunities emerge to combine various sets of satellites to form a temporary constellations to perform collaborative science data collections. Often, new ideas emerge after the launch of satellites. The chance of implementing some of these new mission objectives after launch will depend on if new space assets can be inexpensively and rapidly integrated into temporary or "ad hoc" constellation. On our experiments with the Earth Observing 1 (EO-1) satellite, a New Millennium Program mission, these experiments are being conducted to demonstrate various aspects of an architecture that when taken as a whole will enable progressive mission autonomy and thereby enable the easy cobbling together of heterogeneous satellites via flight and ground software working in concert with ground "wireless access" points for low earth orbiting satellites. The end goal is to provide internet like connectivity to all of the space assets with an interoperable type of architecture. This would enable more science yield from available on orbit instrument resources by managing those resources more wisely. For example, of the 20,000 plus images taken thus far on the EO-1 mission, approximately two thirds are cloudy. Experiments conducted with EO-1 use GOES and other satellite observations of clouds in conjunction with the on-board autonomous rescheduling capability to make real time decisions for which images to take. Increasing the "cloud-free" image yield by just 10-20% could in savings of millions of dollars to the missions

↑
result

since thousands of low value, cloud obscured targets would not be taken and instead would have been substituted for higher value saleable images. Interestingly, we connected hourly GOES observations to task planning for EO-1, so in effect connected GOES and EO-1 as a temporary or "ad hoc" constellation for negligible cost. More details on this particular experiment will be provided later in this paper.

One of the keys to make this concept highly desirable is the interconnectivity for the satellites via the ground. Presently our research is being conducted using existing communications infrastructure which is cumbersome to use at best. Our primary method of connectivity to EO-1 at present is via 11 meter dishes. However, with the emergence of digital signal processing technology, new antenna systems can be created that will be relatively inexpensive and maintenance-free. They would resemble a system of wireless access points rather than the traditional multimillion dollar mechanically driven antenna systems in use today by NASA. Once developed, numerous antenna systems such as these could be deployed around the globe cost-effectively thus providing much more coverage, if not total coverage. Furthermore, as larger amounts of data are required to be downlinked, many lower bandwidth downlink pipes could be substituted thus providing more cost effective space to ground links for large volumes of data. Finally, by connecting these less expensive antennae to the Internet and using protocol standards such as Internet Protocol (IP) and other higher level messaging protocol, easy satellite interoperability can be achieved.

Figure 1 depicts the set of related sensor web tasks that attempt to take steps towards this grand vision using EO-1 as the central hub of these experiments.

2. OPERATIONAL CONCEPT

The key elements of the sensor web include autonomous detection of events, autonomous monitoring of detection notifications, autonomous generation of observation requests, and autonomous rescheduling of observations to acquire data of higher temporal, spatial, and spectral resolution.

The sensor web Architecture has the following design objectives:

- Enable autonomous tasking of the EO-1 spacecraft in response to ground-based science events.
- Support a diverse number of sensor sources and science event types.
- Enable reaction observations to be serviced and uploaded promptly in order to maximize the responsiveness of the sensor web.
- Provide for autonomously executing detection algorithms on-board the satellite
- Deliver detection notifications from the on-board algorithms to the ground

- Allow for detailed tracking of changes made by the sensor web to the mission operations schedule.
- Minimize the impact on the EO-1 operations staff and procedures.

The first autonomy experiment on EO-1 used data from Hyperion, a hyperspectral instrument, to demonstrate the ability to automatically discriminate between clouds, ice/snow, and other high reflection land features on-board the satellite. This detection algorithm was developed by Lincoln Laboratory researchers and was integrated and uploaded to the second of EO-1's dual flight processors by flight software engineers at the Goddard Space Flight Center over two years after the launch of EO-1. Other detection algorithms were developed and now there are five different detection algorithms that routinely run on-board EO-1.

observations based on the detections reported and the overall science goals that are input into the system.

Detection information from other remote sensing platforms is obtained through collaborative efforts with instrument data providers that provide triggering data from installed in-situ sensor suites and low/moderate resolution satellites that directly broadcast their data to ground receiving stations around the world. Data from the sensors is automatically posted to internet web sites in an agreed upon format for retrieval by sensor web monitoring software running on the ground. These monitoring software suites poll sensor sites such as the thermal, seismic, and gas monitoring instruments installed at various volcanoes. The monitoring software also polls various data providers for MODIS, GOES, Quikscat, and other satellite data to obtain event information. The detections made from these data are

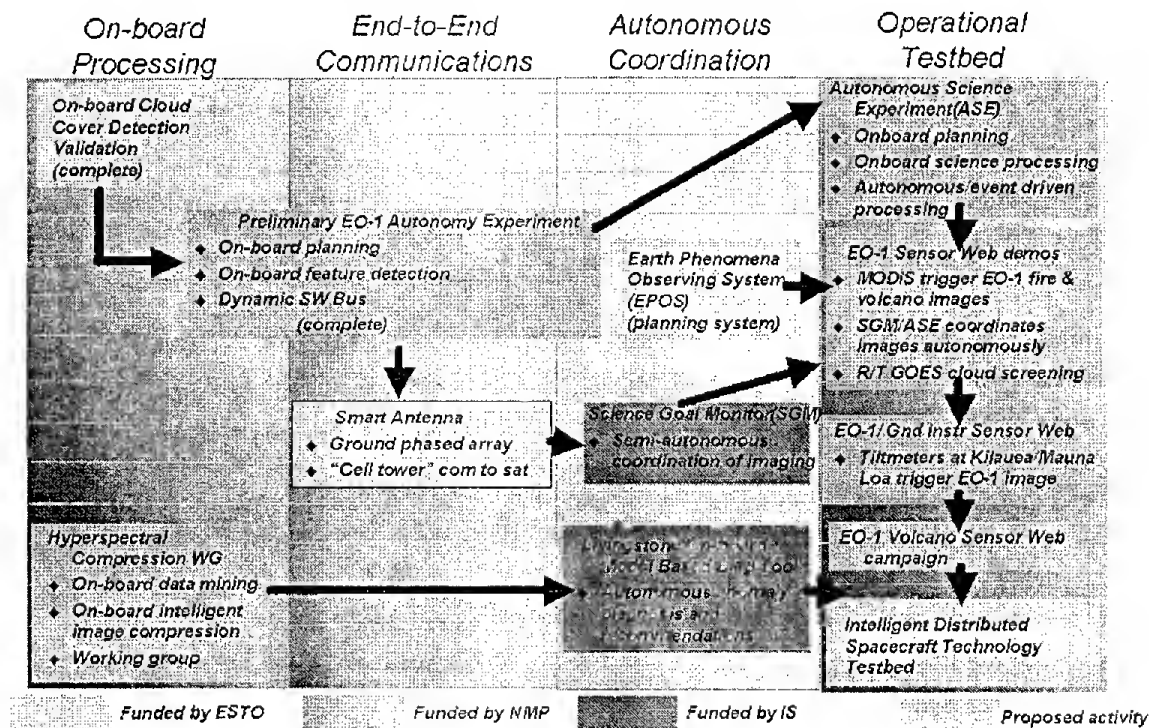


Figure 1 Funded sensor web related experiments using EO-1 as the key platform

These algorithms operate on the hyperspectral data to detect and report on various phenomena within the view of that platform's instruments. Several of these newer algorithms were developed in collaboration with researchers from NASA's Jet Propulsion Laboratory (JPL) in conjunction with the Autonomous Sciencecraft Experiment (ASE). The ASE includes an autonomous on-board scheduling, execution, and replanning algorithm that was integrated and uploaded to control EO-1's observation sequences. The scheduling software responds to the autonomous detections received both from the on-board algorithms as well as those sent from the ground and chains together subsequent EO-1

posted on internet web sites by researchers from agencies such as the USGS, NASA, NRL, and NOAA and universities including the University of Hawaii, University of Arizona, Arizona State, Dartmouth, and other laboratories such as the Draper Lab and Lincoln Lab. The ground monitoring software is integrated with EO-1 satellite operations software at GSFC so that the triggering information can be passed to the spacecraft in time to re-target EO-1 before overflight of the desired target.

On-board EO-1, the scheduling software receives the targeting request, autonomously inserts the observation into the spacecraft schedule, sends commands to

accomplish the maneuver, instrument activation, data recording, and downlink of the data, then activates the on-board autonomous detection algorithms to further examine the new data. New detections made on-board are transmitted to the ground for re-broadcast as triggers for other sensors in the web and are used by the on-board scheduling software to trigger follow-up observations by inserting them into the EO-1 schedule on-board the spacecraft. Future triggered observations are noted on the ground for tracking purposes so that the ground monitoring software is aware of new observations being planned and observations that need to be rescheduled.

The desired operations concept is to extend Internet capabilities to low earth orbiting satellites and to make the space ground interface as seamless and interoperable as possible. Figure 2 depicts how the architecture would look with the addition of antennas that act as wireless access points for various satellites and even other vehicles such as airplanes. Note that in addition to providing continuous connectivity, standards are needed to allow messages to pass seamlessly between software entities in the spacecraft and on the ground. Furthermore, for the long-term vision, software could be uploaded and immediately plugged-in and operate somewhat like JAVA applets.

Note that this architecture would allow space architecture to evolve more rapidly. The traditional approach for building missions involves large-scale system engineering and corresponding costs. In this architecture, many small incremental improvements can be achieved without large expenditures. This has been demonstrated over the last three years as the EO-1 Sensor Web has evolved. Figure 3 depicts part of the connections and triggers established to enable collaborative imaging. One example involving wildfires used MODIS data from Terra and Aqua to detect hot pixels and their locations via the RapidFire workstation in the MODIS instrument center. The Sensor Web monitoring software retrieved the hot pixel locations from RapidFire and sent a trigger to task EO-1 to take a closer look in high resolution via its Advanced Land Imager (ALI) instrument.

3. SENSOR WEB RESEARCH

Experiments to date have focused on the basic capabilities of the various sensor web components. They show the promise of coordinating data from different sources, analyzing the data for a scientifically relevant event, and autonomously updating and rapidly obtaining a follow-on scientifically relevant image in a number of different science domains. The Sensor Web continuously monitors a vast network of in-situ and on-orbit sensors, gathering a holistic view of a science event and tasking the EO-1 spacecraft in response. This process requires a complex combination of ground and on-orbit automation, including coordinating disparate mission planning




Figure 2 Sensor web architecture concept using smart antenna "hot spots" thus enabling extension of the Internet to low earth orbiting satellites. Ultimately this would enable easy interoperability and progressive mission autonomy

systems and functions.

The Sensor Web builds on the successful demonstration of automated mission operations onboard EO-1 using the Autonomous Sciencecraft Experiment (ASE) software system developed by JPL. Baseline EO-1 static observations are selected a week in advance during the Weekly Scene Planning Meeting (WSPM). The observations selected in this meeting are scheduled by the Mission Operations Planning and Scheduling System (MOPSS) each Thursday at the EO-1 Mission Operations Center (MOC). For periods when the spacecraft will be under ASE control, this schedule is then converted to an ASE goal file (INI) and uplinked to EO-1 on Friday for the following week.

ASE. Instead of avoiding contacts with the ground the way the ASE autonomy software was designed, the Sensor Web seeks out contacts to upload last-minute changes in reaction to ground-detected science events. This shift from the current hands-off approach requires a new level of ground-based autonomy to complement the onboard autonomy pioneered by ASE. The two systems must work together in concert to uplink and insert new ground-generated observation requests into the operations schedule.

In order to maximize the responsiveness of the Sensor Web we want to take advantage of as many uplink opportunities as possible—not just the one a week required by ASE. The Sensor Web architecture therefore calls for a procedure to be repeated before each uplink opportunity. This procedure collects




Figure 3 Pictorial representation of some of the sensor web related experiments conducted thus far and also to be conducted

The ASE goal file includes high-level goals for each of the primary spacecraft operations to be performed the following week. These include instrument calibrations, science data collections, and ground contacts. Using an internal model of EO-1, ASE expands these high-level goals into spacecraft activities. At execution time, ASE converts these activities into EO-1 commands, and issues them to the on-board Command and Data Handling (C&DH) subsystem.

The Sensor Web employs a different operations strategy than the one advocated and implemented by

observation requests from the Sensor Web (in response to ground-detected science events), resolves conflicts between new observations and existing observations using observation priority, and uplinks the changes to the onboard schedule to ASE

Ideally the Sensor Web would simply uplink all new observation requests to ASE and would have the onboard planner resolve conflicting observations and generate a new mission plan. Unfortunately the EO-1 spacecraft has limited onboard computational resources (approximately 8 MIPS). Leaving time for on-board planning would require an unacceptable delay between

an incoming observation request and the corresponding overflight opportunity. To work around the performance issue it was decided to maintain a copy of the plan on the ground and compute updates to the onboard plan before each uplink. We thus uplink only changes to the onboard plan, making the Sensor Web the final arbiter of what observations will be collected onboard.

3.1 Software Architecture

The Sensor Web software consists of three primary components that work in concert to recognize science events, generate prioritized observation requests, and insert observations into the EO-1 mission operations schedule (see Figure 4). Science Agents interpret sensor data to extract science phenomena of interest and generate corresponding Science Alerts. The Science Event Manager (SEM) collects Science Alerts, matches them to predefined observation campaigns, and issues prioritized observation requests. The ASPEN planning system inserts these observation requests into the EO-1 mission operations schedule.

3.2 Science Agents

The Science Agents provide the raw data for the Sensor Web. They define the interface between the network of sensors and the SEM. An individual Science Agent may choose to simply pass on data from an in-situ sensor, or synthesize the information from a cluster of sensors, notifying the SEM only on an interpreted science event.

Each Science Alert contains a core set of information required for the SEM to act on the alert. These fields specify the source of the alert, the geographic location of the science phenomenon, and an agent-defined confidence. Additionally a fourth required field classifies the alert within predefined classes of science events.

Agents communicate with the SEM through the XML protocol. No assumptions have been placed on implementation language or location for the individual agents.

Table 1. Science Alert Fields

Required Field	Description
Source	Name and unique ID for source agent—ID assigned by JPL
Target	Geographic location of science alert—specified by a latitude, longitude, and precision
EventTime	Time of the event
ExpirationDate	Time at which the alert is no longer valid
Priority	The agent's perceived priority of the science alert—may be one of LOW, MEDIUM, or HIGH
Event Type	Enumeration of typical science phenomena (VOLCANIC, FLUVIAL, AEOLIAN) and specific subset

3.3 The Science Event Manager

The Science Event Manager (SEM) acts on the Alerts generated by the Science Agents, classifying Alerts by their importance, and generating EO-1 observation requests when Alerts match events specified by participating scientists. These specifications, called Science Campaigns, define the bar by which Alerts become EO-1 observation requests. In some Campaigns a single Alert may generate an observation request, while in others it may take multiple Alerts to generate an observation request. The type or confidence of an alert may affect the priority of the observation request, and with it its likelihood of being scheduled onboard EO-1.

Table 2. Science Campaign Properties

Campaign Field	Description
Title	Name and Unique ID for Campaign
Owner	Requesting Scientist or Institution
Local Priority	LocalPriority: Defines the priority that the scientist sets for a certain event. This priority is used when tasking the asset, but can be superceded by a GlobalPriority (tbd: cleanup)
Global Priority	A priority set by a trusted party that relates this monitor to others monitors in the global context. This priority takes precedences when it conflicts with the LocalPriority set by the scientists. TBD: cleanup, rename
Condition	A specification on a Science Alert or collection of Alerts that when matched generates an observation request.

The core of an observation campaign is the <Condition> block. This block specifies the fields that must be matched in order for an observation request to be generated. The specification takes the form of an evaluation tree, where each condition is logically joined to the next by a Boolean expression. Conditions may use either the EQUALTO, GREATERTHAN, or LESSTHAN operators. More operators may be implemented at a later date.

Each clause within the <Condition> block contains both a <Field> and a <Value>. The <Field> node specifies the value to be extracted from the incoming alert. The <Value> node specifies the absolute value to compare the Alert input against. Additionally the clause may compare two <Field> nodes of an Alert, or compare two <Value> nodes (although this evaluation would always be constant).

For example, the campaign in Figure 4 monitors incoming alerts for those sent from the ModVolc agent. Upon receipt of an Alert, the SEM checks whether the latitude of the target is greater than 34 degrees. If the Alert matches this condition, then the SEM generates an observation request, using the priorities defined by the

rest of the Campaign. The GUI created to facilitate the defining of Campaigns is described in ~~Section 6~~ *Figure 5*. ✓

3.4 ASPEN Mission Planning System

The SEM has the job of recognizing and prioritizing new science events. However the SEM does not have enough information to decide whether the resulting observation request can be achieved within the framework of the current mission plan, or even when the next opportunity to observe the target may take place. These decisions require knowledge of the orbit of the spacecraft, the plan currently executing onboard EO-1, and the constraints on scheduling new observations. Fortunately the core planner of the onboard ASE software, ASPEN, has the capabilities required to make these decisions. The Sensor Web uses ASPEN on the ground to make the above decisions.

Deciding whether an observation can be added to the mission operations schedule presents a number of challenges. First, the observation request must be mapped to a specific target overflight, and new observation goals must be generated. EO-1's orbit nominally allows for up to 10 observations of a target every 16 days (5 daytime and 5 nighttime opportunities), and every observation requires additional goals to point the spacecraft and return it to a nominal state. Second, the new observation will almost always conflict with an existing observation in the onboard schedule. ASPEN must be able to choose between the conflicting observations based on a priority scheme. Third, once a new schedule has been computed, ASPEN must be able to change the onboard schedule to reflect the ground-decisions.

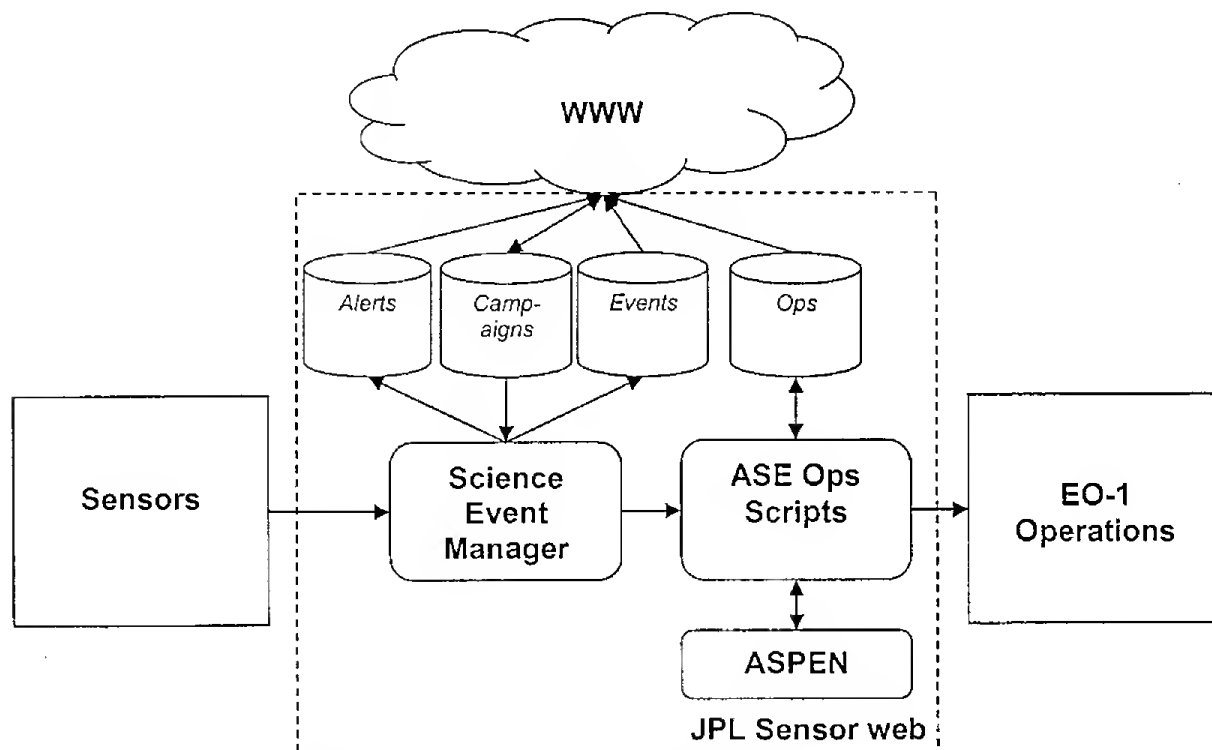


Figure 4 Software Architecture Overview

```

...
<Condition>
  <And>
    <EQUALTO>
      <Field>
        <ScienceEvent>
          <Source/>
        </ScienceEvent>
      </Field>
      <Value>ModVolc</Value>
    </EQUALTO>
    <GREATERTHAN>
      <Field>
        <ScienceEvent>
          <Target>
            <Coordinates>
              <Latitude/>
            </Coordinates>
          </Target>
        </ScienceEvent>
      </Field>
      <Value>34</Value>
    </GREATERTHAN>
  </And>
</Condition>

```

Figure 5 Sample campaign specification

ASPEN generates goals for observation requests using the Momentum Management and Attitude Planning (MMAP) suite of MATLAB tools developed by GSFC. Using MMAP, ASPEN calculates precise overflight times, and determines the parameters for the supporting slews required to point EO-1 at the target. ASPEN performs these calculations as preprocessing before placing an observation request in the proposed schedule.

Once the observation is placed into the schedule, ASPEN resolves any conflicts that may have been created with previous observations. Conflicts are resolved using a priority scheme where every observation is assigned a priority based on EO-1 spacecraft tasking priorities, and scientific value as decided by the SEM.

Finally, with a conflict-free schedule in hand, ASPEN calculates the set of changes required to update the

onboard plan to the new ground-resolved schedule. ASPEN accomplishes this by using a special "user" function that accepts two plans and calculates the goals that have been added or removed from the plan. This function produces a script which can then be uploaded to EO-1 and run within ASE to update the onboard plan. Note that the above scheme requires that the ground schedule is always an exact copy of the onboard schedule. Also, the script described above only operates on top-level goal activities, and not the detailed level of activities that ASPEN uses to plan. The script adds goals to the onboard schedule as "unsatisfied", and then appends commands to have ASE satisfy the goals. After satisfying the goals, ASE will detail the goal into the sub-activities required for execution.

3.5 Operational Constraints

The following sections enumerate and discuss the

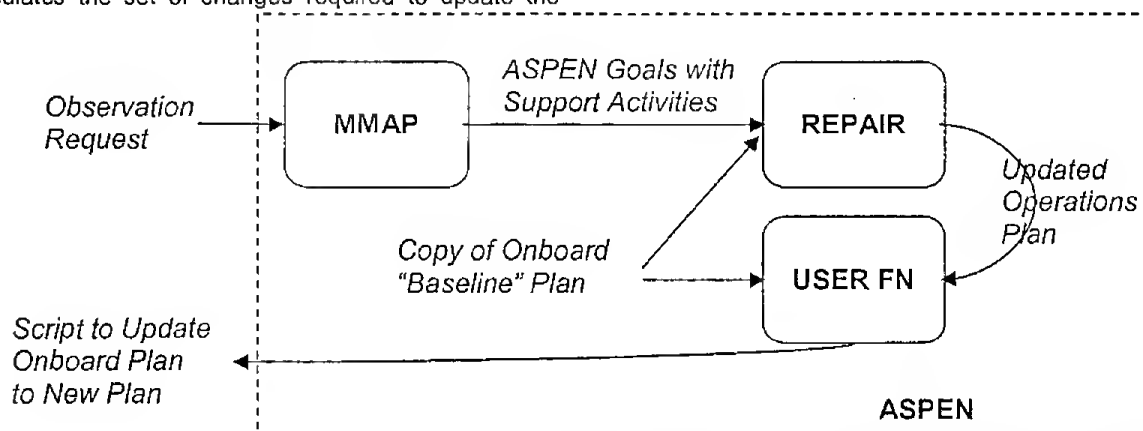


Figure 6 ASPEN information flow

(not mentioned in text)

challenges of integrating the Sensor Web into EO-1 mission operations.

3.6 Modification to EO-1 Planning Process

Integrating Sensor Web observations into the EO-1 planning process requires four critical changes to the EO-1/ASE operations flow described above.

- **Priorities.** Each observation selected in the Weekly Scene Planning Meeting (WSPM) must be assigned a priority based on the scheme outlined in Table 3.
- **MOPSS Report Generation.** The full MOPSS Report for the upcoming week must be generated by the Friday after the WSPM. This report should cover all operations for the following Monday-Sunday.
- **Zero-Bias After Observations.** For the first phase of operations we will need to insert zero-bias activities after each observation. A discussion of this temporary restriction can be found in section 5.3.
- **Notification of Plan Changes.** GSFC must generate MOPSS reports for, and deliver to JPL, any changes to the baseline plan made after the WSPM. See section 5.3 for a more detailed discussion.

Table 3. Observation Priorities

Priority	Description
1	Anomaly Investigation (Spacecraft Health and Welfare)
2	Security (as tasked by NASA HQ)
3	Emergency Response to Natural Disasters or Catastrophic Events (as tasked by NASA HQ)
4	EO-1 Sensor Calibrations and Maintenance
5	Priority Tasked Scenes with Coordinated Ground Truth Measurements (Bulk Customer, Commercial Customer or ALIAS/HIAS with Ground Truth)
6	Bulk Customer or Commercial Customer Paid Scenes, Rotation between groups in highly tasked areas, Otherwise, it is first come, first served.
7	Speculatively Tasked Emergency Response to Natural Disasters or Catastrophic Events.
8	Speculative USGS or NASA Science Collects (internal or external projects), Landsat-7 Islands. Rotation between groups in highly tasked areas.
9	USGS Speculative Collects (Filler scenes defined by Acquisition Coordinator—Driven by gaps in archive)

Note, that we do not require all possible Sensor Web scenarios to be enumerated in the WSPM as was required by our early experiments. In fact, the WSPM does not need to consider the Sensor Web at all outside of assigning priorities to observations. This dramatically reduces the amount of work required in the WSPM to support the Sensor Web.

3.7 EDC Notification and Data Processing

Removing Sensor Web observations from the WSPM allows for increased flexibility in observation campaigns, but introduces complications in the accounting of the observations ASE scheduled and collected. The tracking problem is further compounded by the Sensor Web being event-driven rather than target-driven—targets do not need to be enumerated in a database in order for observations to be scheduled. Observation requests may be generated on the fly for new targets based solely on the latitude and longitude of the detected event.

As the Sensor Web decisions do not involve EDC or EO-1 operations, beyond the assigning of priorities to scenes, record-keeping measures have been implemented to inform interested parties when preemptions occur, including enumerating scenes that were added or removed from the onboard schedule. The following two reports will be delivered to document Sensor Web actions:

- An addendum to the GSFC Daily Report containing the records added and deleted from the baseline schedule delivered by 10:00 GMT the day after a Sensor Web preemption.
- A daily report of Sensor Web activity delivered to EDC user services in comma delineated LTP format.

Additionally Sensor Web scenes have been assigned a unique range of WARP base-ids (1001-1023) to further differentiate from ATS load (1-511) and ASE (512-1000) scenes.

Note that the Sensor Web generates EO-1 scene-ids for all new observations. These scene-ids follow the established EO-1 naming convention shown in Figure 7.

EO1HPPRRRRYYYYDDDDXXXPL	
EO1	Satellite
H	Hyperion Sensor
PPP	Target WRS path
RRR	Target WRS row
YYYY	Year of acquisition
DDD	Julian day of acquisition
X	Hyperion 0=off; 1=on
X	ALI 0=off; 1=on
X	AC 0=off; 1=on
P	Pointing mode
L	Scene length

Figure 7. EO-1 File Naming Convention

3.8 Momentum Management and Attitude Planning

An observation request from the SEM does not contain enough information for ASPEN to generate a corresponding science goal for planning in the mission operations schedule. For ASPEN to plan an observation request, it must not only be able to schedule the request during an overflight of the target, but it must also plan the activities that will point the spacecraft at the target, and ensure that the reaction wheels have the capacity to hold the pointing for the duration of the image. These supporting slew and wheel bias activities have a set of parameters that are calculated through an orbital analysis of the EO-1 spacecraft. A list of the parameters can be found in Table 4.

Table 4. MMAP Parameters

Parameter	
Image Start Time	Slew Quaternion
Image End Time	Slew Wheel Rates
ALI Frame Rate	Wheel Bias Start Time
ALI Integration Time	Wheel Bias Duration
Slew Start	Wheel Bias Rates
Slew Duration	

Managing these support activities represents a challenging problem for the Sensor Web. As each support activity depends on the current momentum of the spacecraft's reaction wheels, and different pointings change this momentum in different ways, we can no longer consider each observation independently. In other words, we can no longer swap two observations without updating the support activities for the observations that follow.

Luckily we can side step this problem by resetting the spacecraft to an initial state after each observation. Resetting or zero-biasing the reaction wheels can accomplish this. Doing so decouples observations, and enables us to do the quick swap that we wish to do. However, the EO-1 operations team does not currently zero-bias after each observation, and we therefore need to make special accommodations during Sensor Web observation windows.

We hope to relax the zero-bias constraint at a later date by tracking the current pointing and momentum states within ASPEN.

3.9 Automated Uploads

In order for the Sensor Web to function autonomously we must also automate the upload of the ground-based Sensor Web schedule changes to the onboard ASE planner. The Sensor Web hooks into the automated commanding capabilities at EO-1 operations created for the Science Goal Monitor (SGM) Gateway scripts.

3.10 Overflight Updates

As with any low earth orbiting satellite, orbit predictions for EO-1 are highly uncertain in the direction of the

satellite's path due to fluctuations in atmospheric drag. Over the course of a week, this error may range from a fraction of a second up to four or more seconds.

Up to this point, the EO-1 operations team has attempted to minimize the impact of this error by delaying command generation to the day before a scheduled observation, thus allowing the use of a one-day predicted spacecraft ephemeris (except for the Monday command load which uses a three-day predict). The integration of the base one-week ASE system followed a similar paradigm, where goals were uploaded a day in advance using the latest ephemeris.

The Sensor Web's automated upload capability allow us to relax this constraint by updating the start time of science goals when new ephemeris information has been received on the ground. These updates are queued along with changes to the onboard plan, allowing for the MOPSS-generated goals to be uploaded only once a week.

3.11 Request Servicing Overhead

The responsiveness of the Sensor Web is bounded by two factors—uplink opportunities and onboard planning time. As such we place two limits on when changes may be made by the Sensor Web:

- All changes must be received 60-min before the uplink.
- Only activities schedule to execute later than 50 min after the uplink may be changed.

The first constraint gives ASPEN on the ground time to generate the changes to the onboard plan. The second constraint derives from the time required onboard to detail the plan (~30 min) and the onboard commit window (~20 min).

3.12 Modification to ASE Operations

Initially no onboard science will be permitted during Sensor Web windows. Reactions to onboard science analysis cause the onboard plan to diverge from the mission operations plan on the ground. As the Sensor Web requires a consistent copy of the onboard plan, and we currently have no mechanism to inform the Sensor Web of the onboard plan changes, we must limit ASE operations to pre-planned scenes only.

3.13 Uplinks/Downlinks

The first release of the Sensor Web software assumes that all uplinks and downlinks are fixed within the operations schedule. This limits the ways in which the Sensor Web may modify the onboard plans, but allows for a greater degree of observability. While scenarios exist where a high priority Sensor Web observation could conflict with an existing downlink, we consider this possibility unlikely and do not expect to relax the fixed contact constraint in the near future.

However, future sensor webs will have continuous access to the ground via smart antenna systems which

emulate wireless access points. The research that we are conducting in line with this vision to enable smart antennas² seeks to eliminate or minimize moving parts thus significantly reducing the cost to purchase and maintain these antenna systems. This in turn will allow the cost-effective deployment of larger networks of antenna systems to radically increase coverage of low earth orbiting satellites. Using the emerging technology of digital signal processing, software is used to shape the antenna pattern. In essence, if taken to the end goal, the software would shape the antenna pattern to follow the target satellite without moving parts such as the large motors used to slew the 11 meter dishes at the NASA Ground Network ground stations. Furthermore, the software would be able to shape the antenna pattern to optimize the desired signal and minimize the impact of interference. Thus whereas most antenna systems have their desired signal diminished by multipath, whereby the same signal bounces off of buildings and other structures to interfere with itself, through the use of this smart antenna technology, the multipath signal can be used to actually enhance the desired signal. Therefore this technology is the perfect technology to create wireless access points for low earth orbiting satellites especially if in the future medium to large constellations are launched. This technology will provide a more cost effective means for a ground station to handle multiple satellites simultaneously. NASA's GN ground stations can typically only handle one satellite per ground station and thus act as bottleneck to potential future constellations. Our research is a collaborative effort between NASA GSFC, NASA Glenn Research Center (GRC), Georgia Institute of Technology and University of Colorado. GSFC provides operational and systems expertise, Glenn and University of Colorado provide miniature phased array expertise and Georgia Institute of Technology provides the adaptive array algorithm expertise and system integration expertise. The effort is a three year effort and has evolved since the inception of the task which began April 2003. At present, the planned milestones along with those already accomplished are as follows:

- Demonstrate a 4 element ground adaptive array in S-band (data rate was 2 kbps) that is able to capture data from EO-1 with no steering. This was successfully completed in April 2004. Figure 7 is a picture of the test setup. Note that the front end elements were comprised of very cheap components such as PVC pipe and wire.
- Demonstrate by April 2005 a 2-4 element adaptive array in X-band (data rate 6 Mbps) using the SAC-C with mechanical steering.
- Demonstrate a 2-4 element adaptive array in X-band (data rate 6 Mbps) using SAC-C with electronic steering.

Figure 8 depicts what a future ground station using this technology to enable sensor webs would look like. It uses a tennis court to depict the size of the system if used for an X-Band link.




Figure 7 Test setup at Georgia Tech for the S-Band test with 4 element adaptive array used with EO-1




Figure 8 Array of electronically steered space fed lens. The outputs of each space fed lens is adaptively combined.

4. AD HOC CONSTELLATION AND SENSOR WEB SCENARIOS

Over the past 3 years a number of scenarios have been executed using a variety of satellites, but with EO-1 as the key satellite since that is the satellite under the control of the sensor web team. The first time that we connected EO-1 to other satellite was in August 2003. At that time, we used the data that Terra and Aqua generate in near realtime via the Rapid Fire workstation in the MODIS instrument center. The Rapid Fire workstation generates hot pixel alerts. Each of the pixels is a 1 km square. During that summer, there were number of large wild fires which are tracked by the Forestry service. Once a fire was identified, MODIS was used to locate hot pixels, which located where the selected fire was burning at present. The Science Goal Monitor (SGM)³ which is located in the EO-1 MOC along with the other planning components automatically generated tasking commands to upload to EO-1 to take a closer look with the Advance Land Imager on EO-1 which is another of the instruments on EO-1. Once the image was taken, the image was transferred to the

Acc
data
→

USGS EROS Data Center for level 0 and level 1 processing. Next the image went to the University of Maryland for synchronizing the image with a map. The image was then transferred to the Forestry Service who used the image to create Burn Area Reflectance Classification (BARC) maps such as the one depicted in Figure 18.




Figure 10 Sample hot pixel map from Rapid Fire workstation in the MODIS instrument center.

It turned out that the Forest Service required a 24–48 hour turn-around for receiving these processed images from the time of the tasking request in order to create the BARC maps in sufficient time to help the BARC team. Although this timeframe is relatively slow, it was fast enough for the Forestry Service for the task at hand and much faster than previously available. The demonstration was done without any of the new antenna technology, however, in the future, the turn-around time for this scenario will improve as new autonomy technology and communications technology are deployed.

Similar scenarios have been deployed using the MODIS instrument to both detect and then take high resolution images of volcanoes through a similar automated ground sequence. In fact, at present, EO-1 is conducting a 500 image volcano campaign which is a combination of triggers from University of Hawaii's MODVOLC website which processes MODIS data and a tilt meter ground instrument trigger at the USGS Hawaiian Volcano Observatory. Furthermore, the Autonomous Sciencecraft Experiment (ASE)⁴ has been demonstrating onboard thermal classification to detect volcanoes and autonomously reschedule EO-1 to image a detected volcano again. The following is a list of the onboard classifiers included with ASE software onboard EO-1 (including the thermal classifier):

- Thermal anomaly detection—uses infrared spectra

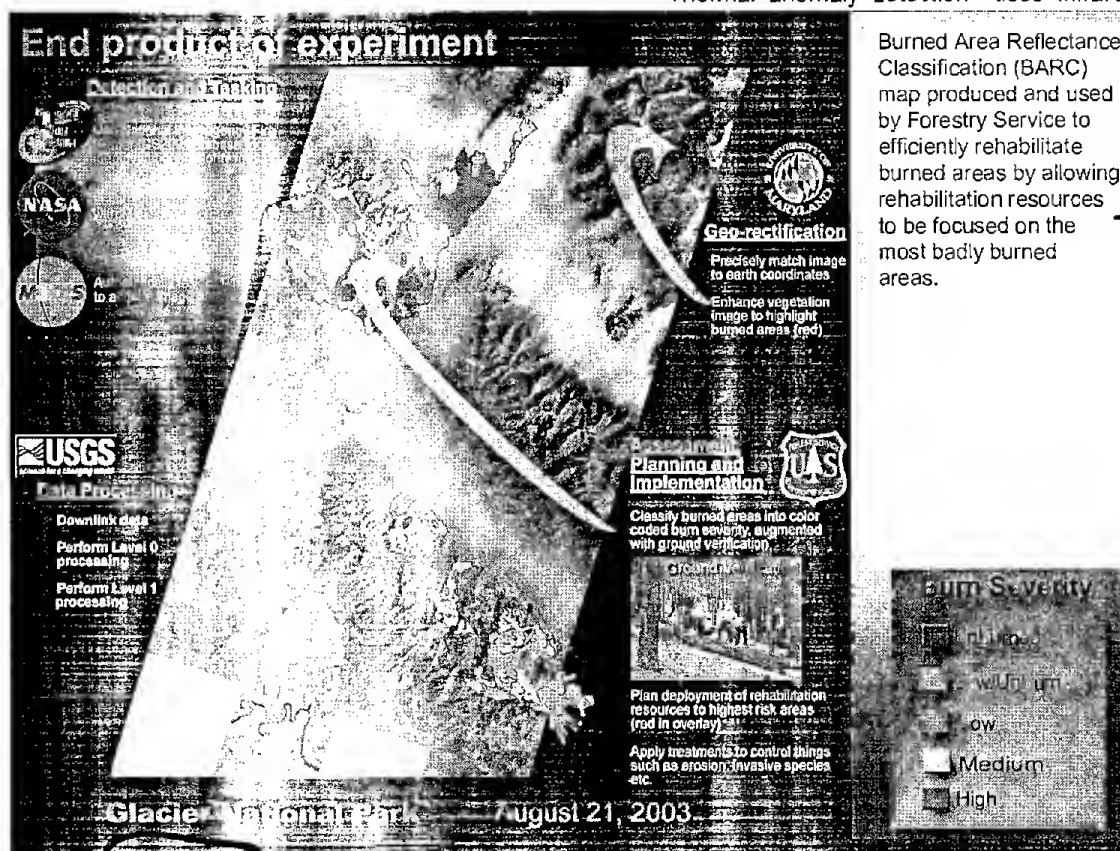


Figure 9 BARC Map created and used by the Forestry Service as a result of the EO-1 Fire Sensor Web experiment. This map was supplied by Rob Sohlberg from the University of Maryland.

Burned Area Reflectance Classification (BARC) map produced and used by Forestry Service to efficiently rehabilitate burned areas by allowing rehabilitation resources to be focused on the most badly burned areas.

Combine this info. with fig 10 caption.

The following types of prototype demonstrations have been conducted on an ongoing basis and are being used to validate the overall concept, calibrate the detection algorithms, and streamline the sensor web interfaces.

4.1 Forest Fires

A number of space based assets are currently being used for detecting active fires, mapping burned area, assessing fire susceptibility and estimating fire emissions. Also when burned areas are re-observed

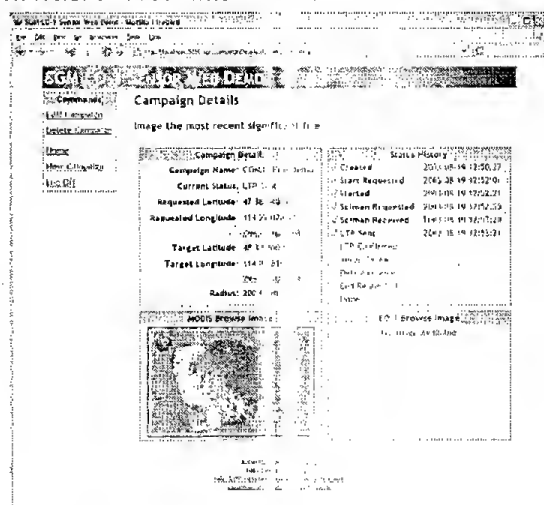


Figure 11 Sample Science Goal Monitor display for transforming high level science goals into specific mission activities

over time, they can help ecologists monitor the recovery of the area. In this domain timely delivery of meteorological and satellite data at the appropriate spatial scale is essential for any of the activities. The sensor web concept is ideally suited for fire management.

Often high-resolution data are needed to provide information at a finer spatial scale, for example to assess fire damage and monitor post fire recovery. To support the needs of the wildfire community, the wildfire sensor web demonstration uses the MODIS RapidFire

data to determine the largest significant fire in a certain region, and then direct EO-1 for a higher resolution image. The significant steps for this scenario are:

- The scientist currently enters their region of interest by entering a latitude, longitude, and radius.
- Sensor web software monitors the daily list of active priority fires from the Remote Sensing Applications Center in Utah (<http://www.fs.fed.us/eng/rsac>), and identifies priority fires from the RSAC that are located within the region of interest.
- Software analyzes the recent history of the fires from the MODIS Rapid Fire data in the region of interest to isolate the latest center of activity.
- A "centroid" of the fire is calculated and coordinates are supplied to the EO-1 planning systems to request a high-priority high-resolution image of the fire and monitor the status of the request.
- A web-based user interface provides the user with a live display of the status of the request and automatically links to the new EO-1 image when it becomes available.


The EO-1 team has performed several demonstrations over the past year, utilizing fires in the southern hemisphere and early spring events in Florida, Southern California and the Southwest U.S.

4.2 Volcanoes

Over 100 volcanoes erupt somewhere on earth every year, sometimes with devastating consequences. Currently, only a limited number of volcanoes are being monitored. When there is a significant event, manual intervention to acquire relevant data is the standard observing strategy. An observing strategy where scientists can “catch” the eruption is extremely useful. A sensor web can provide the means to effectively and efficiently monitor a large number of volcanoes and then obtain timely data to observe an eruption.

In the volcano scenario we focused on demonstrating the ability to conduct long term monitoring of multiple targets, and to conduct follow-on observations when an eruption is detected. The volcano demonstrations show the feasibility of quickly gathering high-resolution satellite images when an event is detected, across federal departments with little human involvement. The significant steps for this scenario are:

- Scientist specifies a prioritized list of volcanoes to monitor for new eruptions.
- Sensor web software monitors each volcano site again using datasets from the MODIS instruments provided by the MODVOLC product at the Hawaii Volcano Observatory.
- When "hot pixels" above the scientist's specified temperature threshold are detected near a volcano site, sensor web software coordinates with EO-1 to automatically request a high-resolution image of the volcano area using the new coordinates. If more than one eruption is detected, the highest priority site



Assets used:

- EO-1
- SPOT
- Aqua & Terra (MODIS)
- Terra (ASTER)
- Landsat 5
- MASTER
 - Aircraft (ER-2) based MODIS & ASTER
- AirDas
 - Airborne Infrared Disaster Assessment System

Figure 12 Sample of fusing observation assets to provide a more complete picture of large LA fires in Nov 2004.

as specified by the science plan is targeted for the follow-on EO-1 observation.

In an extension of the basic volcano demonstration, data from a set of in-situ tilt meter sensors located on Kilauea and Mauna Loa are retrieved. Sensor web software reacts when either tilt meter installation indicates a seismic event or a temperature threshold from MODVOLC data is encountered.

4.3 Floods

Floods affect large areas of the Earth. They are not reliably predictable. Hydrological data from in situ sensors is very sparse. Also in-situ sensors cannot provide information on the extent of the flood. Sensor web observational strategies can be used to map, measure and monitor floods. SGM can be used to automate an on-going monitoring program for a flood area, obtaining images just before the flood, during flood, and then monitoring post-flood recovery progress.

To demonstrate the applicability to monitor river flooding, an interface with the Dartmouth Flood Observatory (<http://www.dartmouth.edu/~floods>) was established. The DFO monitors and posts data from various rivers and wetlands around the world using the QuickSCAT scatterometer instrument. The sensor web software monitors this data for flooding alerts concerning a user-specified river (for our demonstration, the Brahmaputra river in India). The software detected an alert and, as in the fire and volcano scenarios, sent a request for a high-resolution image acquisition to EO-1

based on the latitude and longitude specified in the flood alert record. Future flood scenarios will use ground in-situ sensors to predict potential flooding before it occurs, which will drive subsequent EO-1 observations of the target area similar to the volcano scenario.

4.4 Lake Freezing

The University of Wisconsin maintains a series of buoys in Sparkling Lake that measure surface water temperature. The goal of this scenario was to monitor the data from those buoys to determine when the lake's first freezing occurred, then to take an image of the lake area as soon as possible to characterize the lake environment during the time of transition. SGM monitored the buoy readings for several days and triggered an EO-1 observation as soon as the temperature readings showed the lake's surface was beginning to be freeze. Lake freeze and thaw data is important to shipping interests on the lake.

4.5 Integration of weather forecast to prioritize the satellite tasking

Often many of the detectors on Earth observing satellites observe through clouds and ship the data down. Such data are then useless. To maximize the return of only the most useful scientific data, the observation must be cloud-free. With the increase of onboard processing power, an observation can be analyzed on the satellite itself to determine if the observation is cloud-free before it is transmitted back to earth. In the sensor web domain, we can determine

"near real time" if an observation will have a high probability of being cloud-free by accessing real-time GOES satellite cloud top pressure posted on a NOAA website. The website updates the data for cloud cover over the continental U.S. every hour. If the observation has a high probability of being cloud-free, it can be obtained; else the satellite time can be used more effectively for another observation. In this manner the sensor web can be used to pick from among multiple available targets within view of the satellite to obtain the least cloudy image. In this demonstration, a few hours before a group of potential EO-1 images are to be in view, sensor web software accesses and analyzes the latest GOES cloud cover data and autonomously notifies EO-1 of the least cloudy of the alternates.

4.6 Multi-Phenomenon Monitoring

Our latest demonstrations have begun to mix the above scenarios using scientist-specified priorities so that the system can monitor several emerging events, compare the priorities of the requested images, ensure that the highest priority scene is cloud-free, and then direct EO-1 to pick the highest priority scene at the latest possible moment before the overflight. This allows EO-1 to best use its high demand imaging time while minimizing lost imaging time to cloud-covered scenes. This scenario will allow us to understand how to handle competing priorities, in an environment where there are alternative top-level goals.

5. SOME GOALS FOR SENSOR WEBS AND AD HOC CONSTELLATIONS

The ultimate goal for ad hoc constellations and sensor webs is to respond quickly to transient events and to be able to rapidly reconfigure available assets for science goals. The responsiveness of the end-to-end architecture is limited by the flexibility and speed of communications. Just as the Internet changes in nature with upgrades in performance, so will the sensor webs and ad hoc constellations change as performance

increases. At first, we can only have responses end-to-end in the range of hours to a couple of days. Ultimately, when the speed of communications gets faster and there are more flexible assets then the response time will decrease to minutes thus enabling newer capability. For example, our present sensor web can support wildfire rehabilitation effort because only responses in days are needed. When the sensor web is fast enough, perhaps real-time wild fire management may be enabled.

Taking it one step further, a key future goal is to be able to put together data from multiple satellite sources into a composite picture. Figure 14 depicts an attempt to accomplish this goal with present day assets. In this figure, the Forestry Service integrated images from various satellites and air borne assets for the fires in California during November 2003 which were the largest wildfires in state history. This can be non trivial because the data does not match in terms of things such as pixel size and may require some transformation to make it work. By being able to fuse multiple satellite and airplane sources, they were able to be more responsive to rehabilitating fire damage. ✓

6. CONCLUSION

As EO-1 continues to experiment with various sensor web/ad hoc constellation demonstrations in "slow motion", one of the key catalysts for high performance ad hoc constellations of the future will be cost effective flexible communications. Just as personnel computers were slow and clunky at first, and over time evolved from being able to run spreadsheets to running movie editing programs and various high speed networked applications with the evolution of CPU speed, so will the sensor web become more responsive to faster transient events as the communication system allows faster and more agile communications with various satellite. If we can approach total coverage for low earth orbiting satellites, then there will be an explosion of sensor web/ad hoc constellation applications that will emerge.

ACKNOWLEDGMENTS

Portions of this work were performed at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. We would like to acknowledge the important contributions of Daniel Tran, Gregg Rabideau, Rebecca Castano, Nghia Tang, Michael Burl of JPL, Stephen Ungar and Thomas Brakke of GSFC, Lawrence Ong of SSAI, Jerry Hengemihle, Nick Hengemihle, and Bruce Trout of Microtel LLC, Jeff D'Agostino and Kathleen Blackman of the Hammers Corp., Seth Shulman, Joseph Szwaczkowski, and Robert Bote of Honeywell Corp., Jim Van Gaasbeck and Darrell Boyer of Interface Control Systems, Michael Griffin and Hsiao-hua Burke of MIT Lincoln Laboratory, Ronald Greeley, Thomas Doggett, and Kevin Williams of the Arizona State University, Victor Baker and James Dohm of the University of Arizona, and Rob Sohlberg, Chris Justice, and Jessica McCarty of the University of Maryland, and Anarahda Koratkar of University of Maryland at Baltimore County.

REFERENCES

¹Goddard Space Flight Center, EO-1 Mission page: URL: <http://EO-1.gsfc.nasa.gov> [cited 15 September 2004].

²Ingram, M.A.; Romanofsky, R.; Lee, R.; Miranda, F.; Popovic, Z.; Langley, J.; Barott, W.; Ahmed, M.; Mandl, D.; "Optimizing Satellite Communications with Adaptive and Phased Array Antenna" ESTC 2004, Palo Alto, CA., June 2004

³Koratkar, A; Grosvenor, S.; Jung, J.; Hess, M.; Jones, J.; "Autonomous Multi-sensor Coordination: The Science Goal Monitor", SPIE 2004 Remote Sensing of the Atmosphere, Ocean, Environment, and Space; Honolulu, Hawaii, November 2004

⁴Sherwood, R., Chien, S., Davies, A., Mandl, D., Fryes, S., "Realtime Decision Making on EO-1 Using Onboard Science Analysis" SPIE 2004 Remote Sensing of the Atmosphere, Ocean, Environment, and Space; Honolulu, Hawaii, November 2004

⁵Griffin, M., Burke, H., Mandl, D., and Miller, J., "Cloud Cover Detection Algorithm for the EO-1 Hyperion Imagery," Proceedings of the 17th SPIE AeroSense 2003, Orlando, FL, April 2003.

⁶Burl, M. C., et al. "Automated Detection of Craters and Other Geological Features," Proceedings of the International Symposium on Artificial Intelligence Robotics & Automation in Space, Montreal, Canada, June 2001.

⁷Mandl, D., et al, "Linking Satellites Via Earth "Hot Spots" and the Internet to Form Ad Hoc Constellations", SPIE 2004 Remote Sensing of the Atmosphere, Ocean, Environment, and Space; Honolulu, Hawaii, November 2004

⁸Mandl, D.; "Experimenting with Sensor Webs Using Earth Observing 1", IEEE Aerospace Conference, Big Sky, MT, March 2004